

Approximation Resistant Predicates From Pairwise Independence

Per Austrin*
KTH – Royal Institute of Technology
Stockholm, Sweden

Elchanan Mossel†
U.C. Berkeley
USA

6 Dec 2007

Abstract

We study the approximability of predicates on k variables from a domain $[q]$, and give a new sufficient condition for such predicates to be approximation resistant under the Unique Games Conjecture. Specifically, we show that a predicate P is approximation resistant if there exists a balanced pairwise independent distribution over $[q]^k$ whose support is contained in the set of satisfying assignments to P .

Using constructions of pairwise independent distributions this result implies that

- For general $k \geq 3$ and $q \geq 2$, the MAX k -CSP $_q$ problem is UG-hard to approximate within $q^{\lceil \log_2 k+1 \rceil - k} + \epsilon$.
- For $k \geq 3$ and q prime power, the hardness ratio is improved to $kq(q-1)/q^k + \epsilon$.
- For the special case of $q = 2$, i.e., boolean variables, we can sharpen this bound to $(k + \mathcal{O}(k^{0.525}))/2^k + \epsilon$, improving upon the best previous bound of $2k/2^k + \epsilon$ (Samorodnitsky and Trevisan, STOC'06) by essentially a factor 2.
- Finally, for $q = 2$, assuming that the famous Hadamard Conjecture is true, this can be improved even further, and the $\mathcal{O}(k^{0.525})$ term can be replaced by the constant 4.

1 Introduction

In the MAX k -CSP problem, we are given a set of constraints over a set of boolean variables, each constraint being a boolean function acting on at most k of the variables. The objective is to find an assignment to the variables satisfying as many of the constraints as possible. This problem is NP-hard for any $k \geq 2$, and as a consequence, a lot of research has been focused on studying how well the problem can be approximated. We say that a (randomized) algorithm has

*E-mail: austrin@kth.se. Research funded by Swedish Research Council Project Number 50394001.

†E-mail: mossel@stat.berkeley.edu. Research supported by BSF grant 2004105, NSF CAREER award DMS 0548249 and DOD ONR grant N0014-07-1-05-06

approximation ratio α if, for all instances, the algorithm is guaranteed to find an assignment which (in expectation) satisfies at least $\alpha \cdot \text{Opt}$ of the constraints, where Opt is the maximum number of simultaneously satisfied constraints, over any assignment.

A particularly simple approximation algorithm is the algorithm which simply picks a random assignment to the variables. This algorithm has a ratio of $1/2^k$. It was first improved by Trevisan [22] who gave an algorithm with ratio $2/2^k$ for MAX k -CSP. Recently, Hast [8] gave an algorithm with ratio $\Omega(k/(\log k 2^k))$, which was subsequently improved by Charikar et al. [5] who gave an algorithm with approximation ratio $c \cdot k/2^k$, where $c > 0.44$ is an absolute constant.

The PCP Theorem implies that the MAX k -CSP problem is NP-hard to approximate within $1/c^k$ for some constant $c > 1$. Samorodnitsky and Trevisan [20] improved this hardness to $2^{2\sqrt{k}}/2^k$, and this was further improved to $2^{\sqrt{2k}}/2^k$ by Engebretsen and Holmerin [7]. Finally, Samorodnitsky and Trevisan [21] proved that, if the Unique Games Conjecture [12] is true, then the MAX k -CSP problem is hard to approximate within $2k/2^k$. To be more precise, the hardness they obtained was $2^{\lceil \log_2 k+1 \rceil}/2^k$, which is $(k+1)/2^k$ for $k = 2^r - 1$, but can be as large as $2k/2^k$ for general k . Thus, the current gap between hardness and approximability is a small constant factor of $2/0.44$.

For a predicate $P : \{0, 1\}^k \rightarrow \{0, 1\}$, the MAX CSP(P) problem is the special case of MAX k -CSP in which all constraints are of the form $P(l_1, \dots, l_k)$, where each literal l_i is either a variable or a negated variable. For this problem, the random assignment algorithm achieves a ratio of $m/2^k$, where m is the number of satisfying assignments of P . Surprisingly, it turns out that for certain choices of P , this is the best possible algorithm. In a celebrated result, Håstad [10] showed that for $P(x_1, x_2, x_3) = x_1 \oplus x_2 \oplus x_3$, the MAX CSP(P) problem is hard to approximate within $1/2 + \epsilon$.

Predicates P for which it is hard to approximate the MAX CSP(P) problem better than a random assignment, are called *approximation resistant*. A slightly stronger notion is that of *hereditary* approximation resistance – a predicate P is hereditary approximation resistant if all predicates implied by P are approximation resistant. A natural and important question is to understand the structure of approximation resistance. For $k = 2$ and $k = 3$, this question is resolved – predicates on 2 variables are never approximation resistant, and a predicate on 3 variables is approximation resistant if and only if it is implied by an XOR of the three variables [10, 23]. For $k = 4$, Hast [9] managed to classify most of the predicates with respect to approximation resistance, but for this case there does not appear to be as nice a characterization as there is in the case $k = 3$. It turns out that, assuming the Unique Games Conjecture, most predicates are in fact hereditary approximation resistant – as k grows, the fraction of such predicates tend to 1 [11]. Thus, instead of attempting to understand the seemingly complicated structure of approximation resistant predicates, one might try to understand the possibly easier structure of hereditary approximation resistant predicates, as these constitute the vast majority of *all* predicates.

A natural approach for obtaining strong inapproximability for the MAX k -CSP problem is to search for approximation resistant predicates with very few accepting inputs. This is indeed how all mentioned hardness results for MAX k -CSP come about (except the one implied by the PCP Theorem).

It is natural to generalize the MAX k -CSP problem to variables over a

domain of size q , rather than just boolean variables. Without loss of generality we may assume that the domain is $[q]$. We call this the MAX k -CSP $_q$ problem. For MAX k -CSP $_q$, the random assignment gives a $1/q^k$ -approximation, and any $f(k)$ -approximation algorithm for the MAX k -CSP problem gives a $f(k \lceil \log_2 q \rceil)$ -approximation algorithm for the MAX k -CSP $_q$ problem. Thus, Charikar et al.'s algorithm gives a $0.44k \log_2 q / q^k$ -approximation in the case that q is a power of 2. The best previous inapproximability for the MAX k -CSP $_q$ problem is due to Engebretsen [6], who showed that the problem is NP-hard to approximate within $q^{O(\sqrt{k})}/q^k$.

Similarly to $q = 2$, we can define the MAX CSP(P) problem for $P : [q]^k \rightarrow \{0, 1\}$. Here, there are several natural ways of generalizing the notion of a literal. One possible definition is to say that a literal l is of the form $\pi(x_i)$, for some variable x_i and permutation $\pi : [q] \rightarrow [q]$. A stricter definition is to say that a literal is of the form $x_i + a$, where, again, x_i is a variable, and $a \in [q]$ is some constant. In this paper, we use the second, stricter, definition. As this is a special case of the first definition, our hardness results apply also to the first definition.

1.1 Our contributions

Our main result is the following:

Theorem 1.1. *Let $P : [q]^k \rightarrow \{0, 1\}$ be a k -ary predicate over $[q]$, and let μ be a distribution over $[q]^k$ such that*

$$\Pr_{x \in ([q]^k, \mu)} [P(x)] = 1$$

and for all $1 \leq i \neq j \leq k$ and all $a, b \in [q]$, it holds that

$$\Pr_{x \in ([q]^k, \mu)} [x_i = a, x_j = b] = 1/q^2.$$

Then, for any $\epsilon > 0$, the UGC implies that the MAX CSP(P) problem is NP-hard to approximate within

$$\frac{|P^{-1}(1)|}{q^k} + \epsilon,$$

i.e., P is hereditary approximation resistant.

Using constructions of pairwise independent distributions, we obtain the following corollaries:

Theorem 1.2. *For any $k \geq 3$, $q \geq 2$, and $\epsilon > 0$, it is UG-hard to approximate the MAX k -CSP $_q$ problem within*

$$\frac{q^{\lceil \log_2 k + 1 \rceil}}{q^k} + \epsilon < \frac{k^{\log_2 q} \cdot q}{q^k} + \epsilon.$$

In the special case that $k = 2^r - 1$ for some r the hardness ratio improves to

$$\frac{k^{\log_2 q}}{q^k} + \epsilon.$$

This already constitutes a significant improvement upon the $q^{\mathcal{O}(\sqrt{k})}/q^k$ -hardness of Engebretsen, and in the case that q is a prime power we can improve this even further.

Theorem 1.3. *For any $k \geq 3$, $q = p^e$ for some prime p , and $\epsilon > 0$, it is UG-hard to approximate the MAX k -CSP $_q$ problem within*

$$\frac{k(q-1)q}{q^k} + \epsilon.$$

In the special case that $k = (q^r - 1)/(q - 1)$ for some r , the hardness ratio improves to

$$\frac{k(q-1)+1}{q^k} + \epsilon \leq \frac{kq}{q^k} + \epsilon.$$

Neither of these two theorems improve upon the results of [21] for the case of $q = 2$. However, the following theorem does.

Theorem 1.4. *For any $k \geq 3$ and $\epsilon > 0$, it is UG-hard to approximate the MAX k -CSP problem within*

$$\frac{k + \mathcal{O}(k^{0.525})}{2^k} + \epsilon.$$

If the Hadamard Conjecture is true, it is UG-hard to approximate the MAX k -CSP problem within

$$\frac{4\lceil(k+1)/4\rceil}{2^k} + \epsilon \leq \frac{k+4}{2^k} + \epsilon$$

Thus, we improve the hardness of [21] by essentially a factor 2, decreasing the gap to the best algorithm from roughly $2/0.44$ to roughly $1/0.44$.

1.2 Related work

It is interesting to compare our results to the results of Samorodnitsky and Trevisan [21]. Recall that using the Gowers norm, [21] prove that the MAX k -CSP problem has a hardness factor of $2^{\lceil \log_2 k+1 \rceil}/2^k$, which is $(k+1)/2^k$ for $k = 2^r - 1$, but can be as large as $2k/2^k$ for general k .

Our proof uses the same version of the UGC, but the analysis is more direct and more general. The proof of [21] requires us to work specifically with a linearity hyper-graph test for the long codes. For this test, the success probability is shown to be closely related to the Gowers inner product of the long codes. In particular, in the soundness analysis it is shown that if the value of this test is too large, it follows that the Gowers norm is larger than for “random functions”. From this it is shown that at least two of the functions have large influences which in turns allows us to obtain a good solution for the UGC.

Our construction on the other hand allows any pairwise distribution to define a long-code test. Using [16] we show that if a collection of supposed long codes does better than random for this long code test, then at least two of them have large influences.

Our proof has a number of advantages: first it applies to any pairwise independent distribution. This should be compared to [21] that require us to work

specifically with the hyper-graph linearity test. In particular our results allow us to obtain hardness results for $\text{MAX CSP}(P)$ for a wide range of P 's. The results are general enough to accomodate any domain $[q]$ (it is not clear if the results of [21] extend to larger domains), and we are also able to obtain a better hardness factor for most values of k even in the $q = 2$ case.

Also, our proof uses bounds on expectations of products under certain types of correlation, putting it in the same general framework as many other UGC-based hardness results, in particular those for 2-CSPs [13, 14, 2, 3, 18].

Finally, our proof gives parametrized hardness in the following sense. We give a family of hardness assumptions, called the (t, k) -UGC. All of these assumptions follow from the UGC, and in particular the case $t = 2$ is known to be equivalent to the UGC. However, the (t, k) -UGC assumption is weaker for larger values of t . For each value of t our results imply a different hardness of approximation factor. Specifically, if the (t, k) -UGC is true for some $t \geq 3$, then the $\text{MAX } k\text{-CSP}$ problem is NP-hard to approximate within $\mathcal{O}(k^{\lceil t/2 \rceil - 1} / 2^k)$. Thus, even the $(4, k)$ -UGC gives a hardness of $\mathcal{O}(k/2^k)$, and for $t < \sqrt{k}/\log k$, the (t, k) -UGC gives a hardness better than the best unconditional result known [7].

2 Definitions

2.1 Unique Games

We use the following formulation of the Unique Label Cover Problem: given is a k -uniform hypergraph, where for each edge (v_1, \dots, v_k) there are k permutations π_1, \dots, π_k on $[L]$. We say that an edge (v_1, \dots, v_k) with permutations π_1, \dots, π_k is t -wise satisfied by a labelling $\ell : V \rightarrow [L]$ if there are $i_1 < i_2 < \dots < i_t$ such that $\pi_{i_1}(\ell(v_{i_1})) = \pi_{i_2}(\ell(v_{i_2})) = \dots = \pi_{i_t}(\ell(v_{i_t}))$. We say that an edge is completely satisfied by a labelling if it is k -wise satisfied.

We denote by $\text{Opt}_t(X) \in [0, 1]$ the maximum fraction of t -wise satisfied edges, over any labelling. Note that $\text{Opt}_{t+1}(X) \leq \text{Opt}_t(X)$.

The following conjecture is known to follow from the Unique Games Conjecture (see details below).

Conjecture 2.1. *For any $2 \leq t \leq k$, and $\delta > 0$, there exists an $L > 0$ such that it is NP-hard to distinguish between k -ary Unique Label Cover instances X with label set $[L]$ with $\text{Opt}_k(X) \geq 1 - \delta$, and $\text{Opt}_t(X) \leq \delta$.*

For particular values of t and k we will refer to the corresponding special case of the above conjecture as the (t, k) -Unique Games Conjecture (or the (t, k) -UGC).

Khot's original formulation of the Unique Games Conjecture [12] is then exactly the $(2, 2)$ -UGC, and Khot and Regev [15] proved that this conjecture is equivalent to the $(2, k)$ -UGC for all k , which is what Samorodnitsky and Trevisan [21] used to obtain hardness for $\text{MAX } k\text{-CSP}$.

In this paper, we mainly use the $(3, k)$ -UGC to obtain our hardness results. Clearly, since $\text{Opt}_{t+1}(X) \leq \text{Opt}_t(X)$, the (t, k) -UGC implies the $(t+1, k)$ -UGC, so our assumption is implied by the Unique Games Conjecture. But whether the converse holds, or whether there is hope of proving this conjecture (or, say,

the (k, k) -UGC for large k) without proving the Unique Games Conjecture, is not clear, and should be an interesting direction for future research.

2.2 Influences

It is well known (see e.g. [13]) that each function $f : [q]^n \rightarrow \mathbb{R}$ admits a unique *Efron-Stein decomposition*: $f = \sum_{S \subseteq [n]} f_S$ where

- The function f_S depends on $x_S = (x_i : i \in S)$ only.
- For every $S' \not\subseteq S$, and every $y_{S'} \in [q]^{S'}$ it holds that

$$\mathbb{E}[f_S(x_S) | x_{S'} = y_{S'}] = 0.$$

For $m \leq n$ we write $f^{\leq m} = \sum_{S: |S| \leq m} f_S$ for the m -degree expansion of f . We now define the *influence of the i th coordinate on f* , denoted by $\text{Inf}_i(f)$ by

$$\text{Inf}_i(f) = \mathbb{E}_x [\text{Var}_{x_i}[f(x)]]. \quad (1)$$

We define the *m -degree influence of the i th coordinate on f* , denoted by $\text{Inf}_i^{\leq m}(f)$ by $\text{Inf}_i(f^{\leq m})$.

Recall that the influence $\text{Inf}_i(f)$ measures how much the function f depends on the i 'th variable, while the low degree influences $\text{Inf}_i^{\leq m}(f)$ measures this for the low part of the expansion of f . The later quantity is closely related to the influence of f on “slightly noisy inputs”.

An important property of low-degree influences is that

$$\sum_{i=1}^n \text{Inf}_i^{\leq m}(f) \leq m \text{Var}[f],$$

implying that the number of coordinates with large low-degree influence must be small. In particular, if $f : [q]^n \rightarrow [0, 1]$, then the number of coordinates with low-degree influence at least τ is at most τ/m .

2.3 Correlated Probability Spaces

We will be interested in probability distributions supported in $P^{-1}(1) \subseteq [q]^k$. It would be useful to follow [16] and view $[q]^k$ with such probability measure as a collection of k *correlated spaces* corresponding to the k coordinates. We proceed with formal definitions of two and k correlated spaces.

Definition 2.2. Let (Ω, μ) be a probability space over a finite product space $\Omega = \Omega_1 \times \Omega_2$. The *correlation* between Ω_1 and Ω_2 (with respect to μ) is

$$\rho(\Omega_1, \Omega_2; \mu) = \sup\{ \text{Cov}[f_1(x_1)f_2(x_2)] : f_i : \Omega_i \rightarrow \mathbb{R}, \text{Var}[f_i(x_i)] = 1 \},$$

where (x_1, x_2) is drawn from (Ω, μ) .

Definition 2.3. Let (Ω, μ) be a probability space over a finite product space $\prod_{i=1}^k \Omega_i$, and let $\Omega_S = \prod_{i \in S} \Omega_i$. The correlation of $\Omega_1, \dots, \Omega_k$ (with respect to μ) is

$$\rho(\Omega_1, \dots, \Omega_k; \mu) = \max_{1 \leq i \leq k-1} \rho(\Omega_{\{1, \dots, i\}}, \Omega_{\{i+1, \dots, k\}}; \mu)$$

Of particular interest to us is the case where correlated spaces are defined by a measure that is t -wise independent.

Definition 2.4. Let (Ω, μ) be a probability space over a product space $\Omega = \prod_{i=1}^k \Omega_i$. We say that μ is t -wise independent if, for any choice of $i_1 < i_2 < \dots < i_t$ and b_1, \dots, b_t with $b_j \in \Omega_{i_j}$, we have that

$$\Pr_{w \in (\Omega, \mu)}[w_{i_1} = b_1, \dots, w_{i_t} = b_t] = \prod_{j=1}^t \Pr_{w \in (\Omega, \mu)}[w_{i_j} = b_j]$$

We say that (Ω, μ) is *balanced* if for every $i \in [k], b \in \Omega_i$, we have that $\Pr_{w \in (\Omega, \mu)}[w_i = b] = 1/|\Omega_i|$.

The following theorem considers low influence functions that act on correlated spaces where the correlation is given by a t -wise independent probability measure for $t \geq 2$. It shows that in this case, the functions have almost the same distribution as if they were completely independent. Moreover, the result holds even if some of the functions have large influences as long as in each coordinate not more than t functions have large influences.

Theorem 2.5 ([16], Theorem 6.6 and Lemma 6.9). *Let (Ω, μ) be a finite probability space over $\Omega = \prod_{i=1}^k \Omega_i$ with the following properties:*

- (a) μ is t -wise independent.
- (b) For all $i \in [k]$ and $b_i \in \Omega_i$, $\mu_i(b_i) > 0$.
- (c) $\rho(\Omega_1, \dots, \Omega_k; \mu) < 1$.

Then for all $\epsilon > 0$ there exists a $\tau > 0$ and $d > 0$ such that the following holds. Let f_1, \dots, f_k be functions $f_i : \Omega_i^n \rightarrow [0, 1]$ satisfying that, for all $1 \leq j \leq n$,

$$|\{i : \text{Inf}_j^{\leq d}(f_i) \geq \tau\}| \leq t.$$

Then

$$\left| \mathbb{E}_{w_1, \dots, w_n} \left[\prod_{i=1}^k f_i(w_{1,i}, \dots, w_{n,i}) \right] - \prod_{i=1}^k \mathbb{E}_{w_1, \dots, w_n} [f_i(w_{1,i}, \dots, w_{n,i})] \right| \leq \epsilon,$$

where w_1, \dots, w_n are drawn independently from (Ω, μ) , and $w_{i,j} \in \Omega_i$ denotes the j th coordinate of w_i .

Note that a sufficient condition for (c) to hold in the above theorem is that for all $w \in \Omega$, $\mu(w) > 0$.

Roughly speaking, the basic idea behind the theorem and its proof is that low influence functions cannot detect dependencies of high order – in particular if the underlying measure is pairwise independent, then low influence functions of different coordinates are essentially independent.

3 Main theorem

In this section, we prove our main theorem. Note that it is a generalization of Theorem 1.1.

Theorem 3.1. *Let $P : [q]^k \rightarrow \{0, 1\}$ be a k -ary predicate over a (finite) domain of size q , and let μ be a balanced t -wise independent distribution over $[q]^k$ such that $\Pr_{x \in ([q]^k, \mu)}[P(x)] > 0$. Then, for any $\epsilon > 0$, the $(t+1, k)$ -UGC implies that the MAX CSP(P) problem is NP-hard to approximate within*

$$\frac{|P^{-1}(1)|}{q^k \cdot \Pr_{x \in ([q]^k, \mu)}[P(x)]} + \epsilon$$

In particular, note that if $\Pr_{x \in ([q]^k, \mu)}[P(x)] = 1$, i.e., if the support of μ is entirely contained in the set of satisfying assignments to P , then P is approximation resistant. It is also hereditary approximation resistant, since the support of μ will still be contained in $P^{-1}(1)$ when we add more satisfying assignments to P .

Reduction. Given a k -ary Unique Label Cover instance X , the prover writes down the table of a function $f_v : [q]^L \rightarrow [q]$ for each v , which is supposed to be the long code of the label of the vertex v . Furthermore, we will assume that f_v is folded, i.e., that for every $x \in [q]^k$ and $a \in [q]$, we have $f_v(x + (a, \dots, a)) = f_v(x) + a$ (where the definition of “+” in $[q]$ is arbitrary as long as $([q], +)$ is an Abelian group). When reading the value of $f_v(x_1, \dots, x_L)$, the verifier can enforce this condition by instead querying $f_v(x_1 - x_1, x_2 - x_1, \dots, x_L - x_1)$ and adding x_1 to the result. Let $\eta > 0$ be a parameter, the value of which will be determined later, and define a probability distribution μ' on $[q]^k$ by

$$\mu'(w) = (1 - \eta) \cdot \mu(w) + \eta \cdot \mu_U(w),$$

where μ_U is the uniform distribution on $[q]^k$, i.e., $\mu_U(w) = 1/q^k$. Given a proof $\Sigma = \{f_v\}_{v \in V}$ of supposed long codes for a good labelling of X , the verifier checks Σ as follows.

Algorithm 1: The verifier \mathcal{V}

$\mathcal{V}(X, \Sigma = \{f_v\}_{v \in V})$

- (1) Pick a random edge $e = (v_1, \dots, v_k)$ with permutations π_1, \dots, π_k .
- (2) For each $i \in [L]$, draw w_i randomly from $([q]^k, \mu')$.
- (3) For each $j \in [k]$, let $x_j = w_{1,j} \dots w_{L,j}$, and let $b_j = f_{v_j} \pi_j(x_j)$.
- (4) Accept if $P(b_1, \dots, b_k)$.

Lemma 3.2 (Completeness). *For any δ , if $\text{Opt}_k(X) \geq 1 - \delta$, then there is a proof Σ such that*

$$\Pr[\mathcal{V}(X, \Sigma) \text{ accepts}] \geq (1 - \delta)(1 - \eta) \Pr_{w \in ([q]^k, \mu)}[P(w)]$$

Proof. Take a labelling ℓ for X such that a fraction $\geq 1 - \delta$ of the edges are k -wise satisfied, and let $f_v : [q]^L \rightarrow [q]$ be the long code of the label $\ell(v)$ of vertex v .

Let (v_1, \dots, v_k) be an edge that is k -wise satisfied by ℓ . Then $f_{v_1}\pi_1 = f_{v_2}\pi_2 = \dots = f_{v_k}\pi_k$, each being the long code of $i := \pi_1(\ell(v_1))$. The probability that \mathcal{V} accepts is then exactly the probability that $P(w_i)$ is true, which, since w_i is drawn from $([q]^k, \mu)$ with probability $1 - \eta$, is at least $(1 - \eta) \Pr_{w \in ([q]^k, \mu)}[P(w)]$.

The probability that the edge e chosen by the verifier in step 1 is satisfied by ℓ is at least $1 - \delta$, and so we end up with the desired inequality. \square

Lemma 3.3 (Soundness). *For any $\epsilon > 0$, $\eta > 0$, there is a constant $\delta := \delta(\epsilon, \eta, t, k, q) > 0$, such that if $\text{Opt}_{t+1}(X) < \delta$, then for any proof Σ , we have*

$$\Pr[\mathcal{V}(X, \Sigma) \text{ accepts}] \leq \frac{|P^{-1}(1)|}{q^k} + \epsilon$$

Proof. Assume that

$$\Pr[\mathcal{V}(X, \Sigma) \text{ accepts}] > \frac{|P^{-1}(1)|}{q^k} + \epsilon. \quad (2)$$

We need to prove that this implies that there is a $\delta := \delta(\epsilon, \eta, t, k, q) > 0$ such that $\text{Opt}_{t+1}(X) \geq \delta$.

Equation 2 implies that for a fraction of at least $\epsilon/2$ of the edges e , the probability that $\mathcal{V}(X, \Sigma)$ accepts when choosing e is at least $\frac{|P^{-1}(1)|}{q^k} + \epsilon/2$.

Let $e = (v_1, \dots, v_k)$ with permutations π_1, \dots, π_k be such a “good” edge. For $v \in V$ and $a \in [q]$, define $g_{v,a} : [q]^L \rightarrow \{0, 1\}$ by

$$g_{v,a}(x) = \begin{cases} 1 & \text{if } f_v(x) = a \\ 0 & \text{otherwise} \end{cases}.$$

The probability that \mathcal{V} accepts when choosing e is then exactly

$$\sum_{x \in P^{-1}(1)} \mathbb{E}_{w_1, \dots, w_L} \left[\prod_{i=1}^k g_{v_i, x_i} \pi_i(w_{1,i}, \dots, w_{L,i}) \right],$$

which, by the choice of e , is greater than $|P^{-1}(1)|/q^k + \epsilon/2$. This implies that there is some $x \in P^{-1}(1)$ such that

$$\begin{aligned} \mathbb{E}_{w_1, \dots, w_L} \left[\prod_{i=1}^k g_{v_i, x_i} \pi_i(w_{1,i}, \dots, w_{L,i}) \right] &> |P^{-1}(1)|/q^k + \epsilon' \\ &= \prod_{i=1}^k \mathbb{E}_{w_1, \dots, w_L} [g_{v_i, x_i} \pi_i(w_{1,i}, \dots, w_{L,i})] + \epsilon', \end{aligned}$$

where $\epsilon' = \epsilon/2/|P^{-1}(1)|$ and the last equality uses that, because f_{v_i} is folded and μ is balanced, we have $\mathbb{E}_{w_1, \dots, w_L} [g_{v_i, x_i} \pi_i(w_{1,i}, \dots, w_{L,i})] = 1/q$.

Note that because both μ and μ_U are t -wise independent, μ' is also t -wise independent. Also, we have that for each $w \in [q]^k$, $\mu'(w) \geq \eta/q^k > 0$, which implies both conditions (b) and (c) of Theorem 2.5. Then, the contrapositive formulation of Theorem 2.5 implies that there is an $i \in [L]$ and at least $t + 1$

indices $J \subseteq [k]$ such that $\inf_{\pi_j^{-1}(i)}^{\leq d} (g_{v_j, x_j}) = \inf_i^{\leq d} (g_{v_j, x_j} \pi_j) \geq \tau$ for all $j \in J$, where τ and d are functions of ϵ, η, t, k , and q .

The process of constructing a good labelling of X from this point is standard. For completeness, we give a proof in the appendix. Specifically, Lemma A.1 gives that $\text{Opt}_{t+1}(X) \geq \epsilon/2 \left(\frac{\tau}{d \cdot q}\right)^{t+1}$, which is a function of ϵ, η, t, k , and q , as desired. \square

It is now straightforward to prove Theorem 3.1.

Proof of Theorem 3.1. Let $c = \Pr_{x \in ([q]^k, \mu)}[P(x)]$, $s = |P^{-1}(1)|/q^k$ and $\eta = \min(1/4, \frac{\epsilon c}{4s})$. Note that since the statement of the Theorem requires $c > 0$ we also have $s > 0$ and $\eta > 0$. Assume that the $(t+1, k)$ -UGC is true, and pick L large enough so that it is NP-hard to distinguish between k -ary Unique Label Cover instances X with $\text{Opt}_{t+1}(X) \leq \delta$ and $\text{Opt}_k(X) \geq 1 - \delta$, where $\delta = \min(\eta, \delta(\epsilon c/4, \eta, t, k, q))$, where $\delta(\dots)$ is the function from Lemma 3.3. By Lemmas 3.2 and 3.3, we then get that it is NP-hard to distinguish between MAX CSP(P) instances with $\text{Opt} \geq (1 - \delta)(1 - \eta)c \geq (1 - 2\eta)c$ and $\text{Opt} \leq s + \epsilon c/4$. In other words, it is NP-hard to approximate the MAX CSP(P) problem within a factor

$$\frac{s + \epsilon c/4}{(1 - 2\eta)c} \leq \frac{s(1 + 4\eta)}{c} + (1 + 4\eta)\epsilon/4 \leq s/c + \epsilon$$

\square

4 Inapproximability for MAX k -CSP $_q$

As a simple corollary to Theorem 3.1, we have:

Corollary 4.1. *Let $t \geq 2$ and let μ be a balanced t -wise independent distribution over $[q]^k$. Then the $(t+1, k)$ -UGC implies that the MAX k -CSP $_q$ problem is NP-hard to approximate within*

$$\frac{|\text{Supp}(\mu)|}{q^k}$$

Thus, we have reduced the problem of obtaining strong inapproximability for MAX k -CSP $_q$ to the problem of finding small t -wise independent distributions. As we are mainly interested in the strongest possible results that can be obtained by this method, our main focus will be on pairwise independence, i.e, $t = 2$. However, let us first mention two simple corollaries for general values of t .

For $q = 2$, it is well-known that the binary BCH code gives a t -wise independent distribution over $\{0, 1\}^k$ with support size $\mathcal{O}(k^{\lceil t/2 \rceil})$ [1]. In other words, the $(t+1, k)$ -UGC implies that the MAX k -CSP problem is NP-hard to approximate within $\mathcal{O}(k^{\lceil t/2 \rceil}/2^k)$. Note in particular that the $(4, k)$ -UGC suffices to get a hardness of $\mathcal{O}(k/2^k)$ for MAX k -CSP, which is tight up to a constant factor.

For q a prime power and large enough so that $q \geq k$, there are t -wise independent distributions over $[q]^k$ with support size q^t based on evaluating a random degree- t polynomial over \mathbb{F}_q . Thus, in this setting, the $(t+1, k)$ -UGC implies a hardness factor of q^{t-k} for the MAX k -CSP $_q$ problem.

In the remainder of this section, we will focus on the details of constructions of pairwise independence, giving hardness for MAX k -CSP $_q$ under the $(3, k)$ -UGC.

4.1 Theorems 1.2 and 1.3

The pairwise independent distributions used to give Theorems 1.2 and 1.3 are both based on the following simple lemma, which is well-known but stated here in a slightly more general form than usual:

Lemma 4.2. *Let R be a finite commutative ring, and let $u, v \in R^n$ be two vectors over R such that $u_i v_j - u_j v_i \in R^*$ for some i, j .¹ Let $X \in R^n$ be a uniformly random vector over R^n and let μ be the probability distribution over R^2 of $(\langle u, X \rangle, \langle v, X \rangle) \in R^2$. Then μ is a balanced pairwise independent distribution.*

Proof. Without loss of generality, assume that $i = 1$ and $j = 2$. It suffices to prove that, for all $(a, b) \in R^2$ and any choice of values of X_3, \dots, X_n , we have

$$\Pr[(\langle u, X \rangle, \langle v, X \rangle) = (a, b) \mid X_3, \dots, X_n] = 1/|R|^2.$$

For this to be true, we need that the system

$$\begin{cases} u_1 X_1 + u_2 X_2 = a' \\ v_1 X_1 + v_2 X_2 = b' \end{cases}$$

has exactly one solution, where $a' = a - \sum_{i=3}^n u_i X_i$ and similarly for b' . This in turn follows directly from the condition on u and v . \square

Consequently, given a set of m vectors in R^n such that any pair of them satisfy the condition of Lemma 4.2, we can construct a pairwise independent distribution over R^m with support size $|R|^n$.

Let us now prove Theorem 1.2.

Proof of Theorem 1.2. Let $r = \lceil \log_2 k + 1 \rceil$. For a nonempty $S \subseteq [r]$, let $u_S \in \mathbb{Z}_q^r$ be the characteristic vector of S , i.e., $u_{S,i} = 1$ if $i \in S$, and 0 otherwise. Then, for any $S \neq T$, the vectors u_S and u_T satisfy the condition of Lemma 4.2, and thus, we have that $(\langle u_S, X \rangle)_{S \subseteq [r]}$ for a uniformly random $X \in \mathbb{Z}_q^r$ induces a balanced pairwise independent distribution over $\mathbb{Z}_q^{2^r-1}$, with support size q^r .

When $k = 2^r - 1$ we get a hardness of $q^{\log_2(k)-k}$, but for general values of k , in particular $k = 2^{r-1}$, we may lose up to a factor q . \square

We remark that for $q = 2$ this construction gives exactly the predicate used by Samorodnitsky and Trevisan [21], giving an inapproximability of $2k/2^k$ for all k , and $(k+1)/2^k$ for all k of the form $2^l - 1$.

Intuitively, it should be clear that when we have more structure on R in Lemma 4.2, we should be able to find a larger collection of vectors where every pair satisfies the “independence condition”. This intuition leads us to Theorem 1.3, dealing with the special case of Theorem 1.2 in which q is a prime power. The construction of Theorem 1.3 is essentially the same as that of [17].

¹ R^* denotes the set of units of R . In the case that R is a field, the condition is equivalent to saying that u and v are linearly independent.

Proof of Theorem 1.3. Let $r = \lceil \log_q(k(q-1)+1) \rceil$, and $n = (q^r - 1)/(q - 1) \geq k$.

Let $\mathbb{P}(\mathbb{F}_q^r)$ be the projective space over \mathbb{F}_q^r , i.e., $\mathbb{P}(\mathbb{F}_q^r) = (\mathbb{F}_q^r \setminus 0)/\sim$. Here \sim is the equivalence relation defined by $(x_1, \dots, x_r) \sim (y_1, \dots, y_r)$ if there exists a $c \in \mathbb{F}_q^*$ such that $x_i = cy_i$ for all i , i.e., if (x_1, \dots, x_r) and (y_1, \dots, y_r) are linearly independent. We then have $|\mathbb{P}(\mathbb{F}_q^r)| = (q^r - 1)/(q - 1) = n$.

Choose n vectors $u_1, \dots, u_n \in \mathbb{F}_q^r$ as representatives from each of the equivalence classes of $\mathbb{P}(\mathbb{F}_q^r)$. Then any pair u_i, u_j satisfy the condition of Lemma 4.2, and as in Theorem 1.2, we get a balanced pairwise independent distribution over \mathbb{F}_q^n , with support size q^r .

When $k = (q^r - 1)/(q - 1)$, this gives a hardness of $k(q - 1) + 1$, and for general k , in particular $k = (q^{r-1} - 1)/(q - 1) + 1$, we lose a factor q in the hardness ratio. \square

Again, for $q = 2$, this construction gives the same predicate used by Samorodnitsky and Trevisan. In the case that $q \geq k$, we get a hardness of q^2/q^k , the same factor as we get from the general construction for t -wise independence mentioned at the beginning of this section.

4.2 Theorem 1.4

Let us now look closer at the special case of boolean variables, i.e., $q = 2$. So far, we have only given a different proof of Samorodnitsky and Trevisan's result, but we will now show how to improve this.

An Hadamard matrix is an $n \times n$ matrix over ± 1 such that $HH^T = nI$, i.e., each pair of rows, and each pair of columns, are orthogonal. Let $h(n)$ denote the smallest $n' \geq n$ such that there exists an $n' \times n'$ Hadamard matrix. It is a well-known fact that Hadamard matrices give small pairwise independent distributions and thus give hardness of approximating MAX k -CSP. To be specific, we have the following proposition:

Proposition 4.3. *For every $k \geq 3$, the $(3, k)$ -UGC implies that the MAX k -CSP problem is UG-hard to approximate within $h(k+1)/2^k + \epsilon$.*

Proof. Let $n = h(k+1)$ and let A be an $n \times n$ Hadamard matrix, normalized so that one column contains only ones. Remove $n - k$ of the columns, including the all-ones column, and let A' be the resulting $n \times k$ matrix. Let $\mu : \{-1, 1\}^k \rightarrow [0, 1]$ be the probability distribution which assigns probability $1/n$ to each row of A' . Then μ is a balanced pairwise independent distribution with $|\text{Supp}(\mu)| = h(k+1)$. \square

It is well known that Hadamard matrices can only exist for $n = 1$, $n = 2$, and $n \equiv 0 \pmod{4}$. The famous *Hadamard Conjecture* asserts that Hadamard matrices exist for all n which are divisible by 4, in other words, that $h(n) = 4\lceil n/4 \rceil \leq n + 3$. It is also possible to get useful unconditional bounds on $h(n)$. We now give one such easy bound.

Theorem 4.4 ([19]). *For every odd prime p and integers $e, f \geq 0$, there exists an $n \times n$ Hadamard matrix H_n where $n = 2^e(p^f + 1)$, whenever this number is divisible by 4.*

Theorem 4.5 ([4]). *There exists an integer n_0 such that for every $n \geq n_0$, there is a prime p between n and $n + n^{0.525}$.*

Corollary 4.6. *We have: $h(n) \leq n + \mathcal{O}(n^{0.525})$.*

Proof. Let p be the smallest prime larger than $n/2$, and let $n' = 2(p+1) \geq n$. Then, Theorem 4.4 asserts that there exists an $n' \times n'$ Hadamard matrix, so $h(n) \leq n'$. If n is sufficiently large ($n \geq 2n_0$), then by Theorem 4.5, $p \leq n/2 + (n/2)^{0.525}$ and $n' \leq n + 2n^{0.525}$, as desired. \square

Theorem 1.4 follows from Proposition 4.3 and Corollary 4.6.

It is probably possible to get a stronger unconditional bound on $h(n)$ than the one given by Corollary 4.6, by using stronger construction techniques than the one of Theorem 4.4.

5 Discussion

We have given a strong sufficient condition for predicates to be hereditary approximation resistant under (a weakened version of) the Unique Games Conjecture: it suffices for the set of satisfying assignments to contain a balanced pairwise independent distribution. Using constructions of small such distributions, we were then able to construct approximation resistant predicates with few accepting inputs, which in turn gave improved hardness for the MAX k -CSP $_q$ problem.

There are several aspects here where there is room for interesting further work:

As mentioned earlier, we do not know whether the (t, k) -UGC implies the “standard” UGC for large values of t . In particular, proving the (t, k) -UGC for some $t < \sqrt{k}/\log k$ would give hardness for MAX k -CSP better than the best current NP-hardness result. But even understanding the (k, k) -UGC seems like an interesting question.

A very natural and interesting question is whether our condition is also necessary for a predicate to be hereditary approximation resistant, i.e., if pairwise independence gives a complete characterization of hereditary approximation resistance.

Finally, it is natural to ask whether our results for MAX k -CSP $_q$ can be pushed a bit further, or whether they are tight. For the case of boolean variables, Hast [9] proved that any predicate accepting at most $2\lfloor k/2 \rfloor + 1$ inputs is *not* approximation resistant. For $k \equiv 2, 3 \pmod{4}$ this exactly matches the result we get under the UGC and the Hadamard Conjecture (which for $k = 2^r - 1$ and $k = 2^r - 2$ is the same hardness as [21]). For $k \equiv 0, 1 \pmod{4}$, we get a gap of 2 between how few satisfying assignments an approximation resistant predicate can and cannot have.

Thus, the hitherto very successful approach of obtaining hardness for MAX k -CSP by finding “small” approximation resistant predicate, can not be taken further, but there is still a small constant gap of roughly $1/0.44$ to the best current algorithm. It would be interesting to know whether the algorithm can be improved, or whether the hardest instances of MAX k -CSP are not MAX CSP(P) instances for some approximation resistant P .

For larger q , this situation becomes a lot worse. When $q = 2^l$ and $k = (q^r - 1)/(q - 1)$, we have a gap of $\Theta(q/\log_2 q)$ between the best algorithm and the best inapproximability, and for general values of q and k , the gap is even larger.

References

- [1] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.
- [2] Per Austrin. Balanced Max 2-Sat Might Not be the Hardest. In *ACM Symposium on Theory of Computing (STOC)*, pages 189–197, 2007.
- [3] Per Austrin. Towards Sharp Inapproximability For Any 2-CSP. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 307–317, 2007.
- [4] R. C. Baker, G. Harman, and J. Pintz. The Difference Between Consecutive Primes, II. *Proceedings of the London Mathematical Society*, 83(3):532–562, 2001.
- [5] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Approximation Algorithm for the Max k-CSP Problem, 2006.
- [6] Lars Engebretsen. The nonapproximability of non-boolean predicates. *SIAM Journal on Discrete Mathematics*, 18(1):114–129, 2004.
- [7] Lars Engebretsen and Jonas Holmerin. More Efficient Queries in PCPs for NP and Improved Approximation Hardness of Maximum CSP. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 194–205, 2005.
- [8] Gustav Hast. Approximating Max kCSP – Outperforming a Random Assignment with Almost a Linear Factor. In *ICALP 2005*, pages 956–968, 2005.
- [9] Gustav Hast. *Beating a Random Assignment – Approximating Constraint Satisfaction Problems*. PhD thesis, KTH – Royal Institute of Technology, 2005.
- [10] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [11] Johan Håstad. On the approximation resistance of a random predicate. To appear in RANDOM-APPROX, 2007.
- [12] Subhash Khot. On the power of unique 2-prover 1-round games. In *STOC 2002*, pages 767–775, 2002.
- [13] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *Siam Journal on Computing*, 37:319–357, 2007.
- [14] Subhash Khot and Ryan O’Donnell. SDP gaps and UGC-hardness for MAXCUTGAIN. In *FOCS 2006*, pages 217–226, 2006.
- [15] Subhash Khot and Oded Regev. Vertex Cover Might be Hard to Approximate to within $2 - \epsilon$. In *IEEE Conference on Computational Complexity*, pages 379–, 2003.

- [16] Elchanan Mossel. Gaussian bounds for noise correlation of functions. arXiv Report math/0703683v3, 2007.
- [17] G. L. O'Brien. Pairwise Independent Random Variables. *Annals of Probability*, 8(1):170–175, 1980.
- [18] Ryan O'Donnell and Yi Wu. An optimal SDP algorithm for Max-Cut, and equally optimal Long Code tests. Manuscript, 2007.
- [19] Raymond E. A. C. Paley. On orthogonal matrices. *Journal of Mathematics and Physics*, 12:311–320, 1933.
- [20] Alex Samorodnitsky and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *STOC*, pages 191–199, 2000.
- [21] Alex Samorodnitsky and Luca Trevisan. Gowers uniformity, influence of variables, and PCPs. In *STOC 2006*, pages 11–20, 2006.
- [22] Luca Trevisan. Parallel Approximation Algorithms by Positive Linear Programming. *Algorithmica*, 21:72–88, 1998.
- [23] Uri Zwick. Approximation Algorithms for Constraint Satisfaction Problems Involving at Most Three Variables Per Constraint. In *SODA 1998*, 1998.

A Good labellings from influential variables

Lemma A.1. *Let X be a k -ary Unique Label Cover instance. Furthermore, for each vertex v , let $f_v : [q]^k \rightarrow [q]$ and define*

$$g_{v,a}(x) = \begin{cases} 1 & \text{if } f_{v_i} = a \\ 0 & \text{otherwise} \end{cases}.$$

Then if there is a fraction of at least ϵ edges $e = (v_1, \dots, v_k)$ with a vector $a \in [q]^k$, an index $i \in [L]$ and a set $J \subseteq [k]$ of $|J| = t$ indices such that

$$\inf_{\pi_j^{-1}(i)}^{\leq d} (g_{v_j, a_j}) \geq \tau \tag{3}$$

for all $j \in J$, then $\text{Opt}_t(X) \geq \delta := \epsilon \left(\frac{\tau}{d \cdot q} \right)^t$.

Proof. For each $v \in V$, let

$$C(v) = \{ i \mid \inf_i^{\leq d} (g_{v,a}) \geq \tau \text{ for some } a \in [q] \}.$$

Note that $|C(v)| \leq q \cdot d / \tau$.

Define a labelling $\ell : V \rightarrow [L]$ by picking, for each $v \in V$, a label $\ell(v)$ uniformly at random from $C(v)$ (or an arbitrary label in case $C(v)$ is empty). Let $e = (v_1, \dots, v_k)$ be an edge satisfying Equation 3. Then for all $j \in J$, $\pi_j^{-1}(i) \in C(v_j)$, and thus, the probability that $\pi_j(\ell(v_j)) = i$ is $1/|C(v_j)|$. This implies that the probability that this edge is t -wise satisfied is at least $\prod_{j \in J} 1/|C(v_j)| \geq \left(\frac{\tau}{d \cdot q} \right)^t$. Overall, the total expected number of edges that are t -wise satisfied by ℓ is at least $\delta = \epsilon \left(\frac{\tau}{d \cdot q} \right)^t$, and thus $\text{Opt}_t(X) \geq \delta$. \square